

This listing of claims will replace all prior versions, and listings, of claims in the application:

Listing of Claims:

1-12. (Cancelled)

13. (Previously Presented) A method for task processing and monitoring of configuration and state information, comprising:

detecting a server comprising two clusters, wherein the server is adapted to perform tasks issued by a browser, wherein each of the two clusters is aware of a first server process and a second server process, wherein the first server process and the second server process are each on a different cluster configured to be a domain server, wherein the first server process and the second server process each have a list of agent processes within a domain that are registered with that server process, and wherein the first server process and the second server process each forward each task issued by the browser to that server process to a registered agent process to perform that task;

at each of the two clusters,

registering a first agent process with a first server process to notify the first server process that the first agent process exists to perform tasks for the first server process to complete the tasks issued by the browser, wherein the first agent process and the first server process form a first agent/server pair;

registering a second agent process with a second server process to notify the second server process that the second agent process exists to perform tasks for the second server process to complete the tasks issued by the browser, wherein the second server process is different from the first server process with which the first agent process is registered, wherein the second agent process and the second server process form a second agent/server pair;

when a task is to be executed by the first server process, executing the task with the first agent process;

when the task is to be executed by the second server process, executing the task with the second agent process; and

wherein when one of the first agent/server pair and the second agent/server pair fails, the other of the first agent/server pair and the second agent/server pair continues processing in the cluster.

14. (Previously Presented) The method of claim 13, further comprising:
at each of the two clusters, storing configuration and state information for one or more storage devices accessed by that cluster as persistent data at the cluster, wherein the configuration information includes how many storage devices are in the cluster, and wherein the state information includes an indication of whether each storage device is available or unavailable.

15. (Previously Presented) The method of claim 14, wherein the first agent process is launched at one of the two clusters and further comprising:
collecting configuration information, including how many storage devices are in the cluster, and state information, including whether each storage device is available or unavailable;
storing the configuration and state information as persistent data at the cluster;
under control of the first agent process in the first agent/server pair,
(i) retrieving the stored configuration and state information; and
(ii) transmitting the retrieved configuration and state information to the first server process in the first agent/server pair.

16. (Previously Presented) The method of claim 13, wherein the second agent process is launched at one of the two clusters and further comprising:
collecting configuration information, including how many storage devices are in the cluster, and state information, including whether each storage device is available or unavailable;
storing the configuration and state information as persistent data at the cluster;
under control of the second agent process in the second agent/server pair,
(i) retrieving the stored configuration and state information; and
(ii) transmitting the retrieved configuration and state information to the second server process in the second agent/server pair.

17. (Previously Presented) The method of claim 13, wherein the first agent process is launched at one of the two clusters and further comprising:

receiving at least one of changed configuration information and changed state information for the cluster, wherein the configuration information includes how many storage devices are in the cluster, and wherein the state information includes an indication of whether each storage device is available or unavailable;

storing the at least one of changed configuration information and changed state information as persistent data at the cluster;

broadcasting the at least one of changed configuration information and changed state information for the cluster; and

under control of the first agent process in the first agent/server pair,

(i) retrieving the stored at least one of changed configuration information and state information; and

(ii) transmitting the retrieved at least one of changed configuration information and state information to the first server process in the first agent/server pair.

18. (Previously Presented) The method of claim 14, wherein the second agent process is launched at one of the two clusters and further comprising:

receiving at least one of changed configuration information and changed state information for the cluster, wherein the configuration information includes how many storage devices are in the cluster, and wherein the state information includes an indication of whether each storage device is available or unavailable;

storing the at least one of changed configuration information and changed state information as persistent data at the cluster;

broadcasting the at least one of changed configuration information and changed state information for the cluster; and

under control of the second agent process in the second agent/server pair,

(i) retrieving the stored at least one of changed configuration information and state information; and

(ii) transmitting the retrieved at least one of changed configuration information and state information to the second server process in the second agent/server pair.

19. (Original) The method of claim 13, wherein the first agent process is launched if a first server is configured and wherein the second agent process is launched if a second server is configured.
20. (Original) The method of claim 13, further comprising:
under control of the first agent process,
receiving a request to execute the task from the first server process;
storing identification for the first agent process in persistent data;
invoking a driver process for executing the task;
receiving task completion status from the driver process; and
forwarding the task completion status to the first server process.
21. (Original) The method of claim 13, further comprising:
under control of the second agent process,
receiving a request to execute the task from the second server process;
storing identification for the second agent process in persistent data;
invoking a driver process for executing the task;
receiving task completion status from the driver process; and
forwarding the task completion status to the second server process.

22-30. (Cancelled)